

Neural Bookmarks: Information Retrieval with Deep Learning and EEG Data

Glenn Bruns, Michael Haidar

California State University, Monterey Bay
Seaside, CA 93955
gbruns@csumb.edu, mhaidar@csumb.edu

Abstract

In neural memory decoding, a concept being mentally recalled is identified using brain data. Recently, the feasibility of neural memory decoding with EEG data has been demonstrated. Here we propose a new application – neural information retrieval – that uses neural memory decoding to allow a document to be retrieved merely by thinking about it. In this paper we describe neural memory decoding, define the application of neural information retrieval, present experimental results related to the practicality of the application, and discuss issues of deployment and data privacy.

Introduction

In the near-future, we will not type commands, or say commands, but will think them. Already, brain-computer interfaces are in use, and it is possible to produce a transcript of what is being silently read from a record of brain activity obtained using functional magnetic resonance imaging (fMRI) (Tang et al. 2023).

Recently, neural memory decoding with electroencephalogram (EEG) imaging has been shown to be possible (Bruns, Haidar, and Rubino 2023). Neural memory decoding is the reconstruction of a mentally-recalled concept from brain data. This finding is important because, unlike fMRI, EEG devices can be inexpensive and comfortable. Consumers could purchase EEG devices and wear them for much of the day. Neural memory decoding with consumer-grade EEG devices suggests a range of exciting applications.

Neural information retrieval is one such application. As more information becomes available on the internet, finding a document that one has encountered earlier becomes a difficult problem. Creating and storing bookmarks is not a good solution, because user-created labels or keywords are subjective and hard to create. The alternative of re-searching for a document can be difficult and time consuming.

With neural information retrieval, information seen once can be retrieved by merely thinking about it. In this application (see Fig. 1) a user, after finding a useful document or website, thinks about it briefly while a short EEG is recorded. Later, to retrieve the document, the user briefly recalls the contents of the document while an EEG is again

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

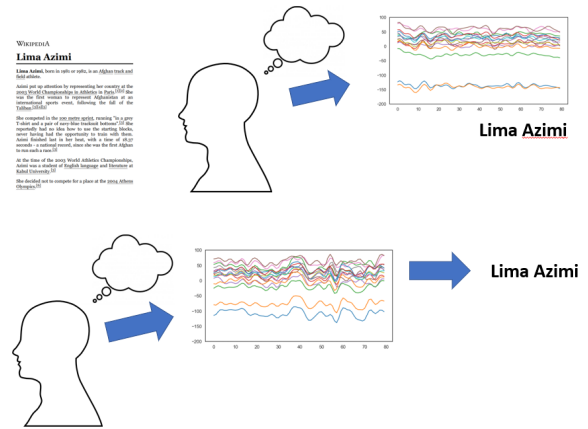


Figure 1: Neural information retrieval. Top: To index, or “bookmark” a document, the system records the user’s brain activity and the document location. Bottom: To search for a document, the system captures brain activity while the user recalls the document. The user is then presented with a list of predicted document links. (Figure from (Bruns, Haidar, and Rubino 2023); used with permission.)

captured. The system then presents links to documents as search results. In contrast to keywords, an EEG provides an information-rich bookmark and does not require creative work by the user.

The model used in (Bruns, Haidar, and Rubino 2023) to show the feasibility of neural memory decoding with EEG data works by mapping segments of an EEG to a lower-dimensional embedding space, classifying each embedding, and then using an ensemble method to yield a class for the EEG as a whole. We will later explain the deployment and data privacy benefits in adopting this model for neural information retrieval.

Our main contributions are as follows: we define the neural information retrieval application, show how the design of (Bruns, Haidar, and Rubino 2023) can be adapted to implement the application, and provide experimental results concerning the practicality of the application.

In what follows we describe the problem of neural memory decoding, briefly survey existing applications of neural

decoding, define the application of neural information retrieval, present experimental results related to the question of whether the application is practical, and discuss deployment and data privacy issues.

Neural Memory Decoding

Neural decoding is the reconstruction of stimuli or mental state from a record of brain activity. For example, neural decoding can be used to reconstruct an individual’s emotional state from an EEG (Bird et al. 2019).

EEG and fMRI are two widely-used methods for measuring brain activity. Some advantages of EEG are that EEG devices are much cheaper, much smaller, and have better temporal resolution. An advantage of fMRI is its high spatial resolution. EEG and fMRI can be used simultaneously (Ritter and Villringer 2006), and other methods also exist to measure brain activity, such as positron emission tomography (PET), and magnetoencephalography (MEG).

Neural memory decoding is the neural decoding of memory. Memory can be categorized as either episodic (memory of events) or semantic (memory of facts, untied to events) (Tulving 1972).

Existing work on the decoding of episodic memory includes classifying whether participants had previously seen pictures of faces (Rissman, Greely, and Wagner 2010), classifying which of three film clips participants had recalled (Chadwick et al. 2010), and detecting deception while participants respond to questions about personal experiences (Ofen et al. 2017). In these works, fMRI was used.

With respect to semantic memory decoding, Polyn *et al.* classify brain activity captured during the study and recall of photographs using fMRI data, distinguishing between the categories of face, location, and object (Polyn et al. 2006). Classifier accuracy is not reported. Abstract concepts such as *multiplication* and *consciousness* have also been decoded using fMRI data (Vargas and Just 2020).

The feasibility of semantic memory decoding using EEG data was demonstrated in (Bruns, Haidar, and Rubino 2023). We now sketch the experimental setup of that work; see (Bruns, Haidar, and Rubino 2023) for details. Data for the study was collected as follows. First, a set of 103 Wikipedia pages were identified and placed into groups of 10. Then, for each group:

- On “day 0”, the study participant read each page in the group for 5 minutes, then mentally recalled each page for 75 seconds while a 16-channel EEG was recorded.
- The next day (“day 1”), the participant again mentally recalled each page for 75 seconds while EEG data was recorded.
- Two days later (“day 3”), the participant repeated the last step.

Data was collected on multiple days because memories of semantic concepts change over time (Winocur and Moscovitch 2011).

After data preprocessing, experiments were then conducted as follows, with the aim of seeing whether a page could be identified from a day 1 EEG given a collection of labeled day 0 EEGs:

- The Wikipedia pages were split randomly into training and test sets.
- A classifier was trained on the day 0 and day 1 EEGs for the training pages, with the page titles (or “topics”) serving as training labels. The classifier was also provided with the day 0 EEGs for the test pages, but no training was performed with this data.
- The classifier was then presented with day 1 EEGs for the test pages, and the label of each EEG was inferred.
- Top-1, top-2, and top-3 accuracy values were computed.

The above steps were repeated, yielding average accuracy values. (This process is sometimes called “repeated random sub-sampling validation”.) Note that the classifier was tested on pages not used in training.

With a test set size of 25 topics, the classifier obtained a mean top-1 accuracy of 75.0%, and a mean top-3 accuracy of 94.4% (a top-1 accuracy of 4% would be obtained by chance). Similar results were obtained when day 3 EEG data was used in place of day 1 data.

We now describe details of the classifier used in (Bruns, Haidar, and Rubino 2023), because it will be adapted for the application we propose. For simplicity, we refer to that system as EMD, for “Experimental Memory Decoder”.

The classifier of EMD is an ensemble that classifies an EEG by first segmenting the EEG, classifying each segment, and then combining the inferred segment labels to yield a label for the EEG as a whole. Fig. 2 illustrates the process. The figure shows only four segments, but in reality an EEG of 75 seconds yields about 850 segments when using a rolling (or “sliding”) window with a length of 100 samples and a stride of 10 samples. We use “segment” to refer to a rolling-window sample of an EEG.

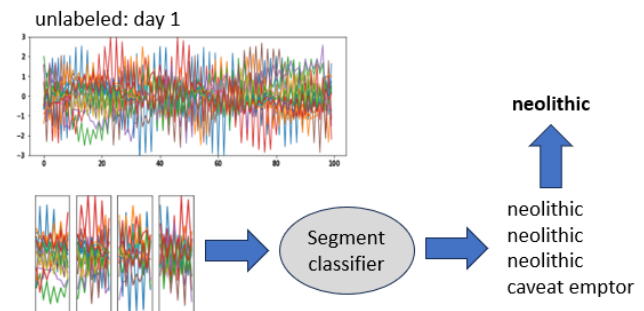


Figure 2: From the inferred labels for every EEG segment, a label for the EEG as a whole is inferred.

Supervised representation learning (Bengio, Courville, and Vincent 2013) is used to classify segments. With representation learning, classification proceeds in two stages. In the “upstream” stage, a convolutional embedding model maps EEG segments to points in an embedding space. The embedding model is trained using supervised contrastive loss (Khosla et al. 2020): the aim is that the embeddings of two segments belonging to the same EEG (in other words, the same topic) will be nearby in the embedding space. In the

“downstream” stage, a k-nearest-neighbor (KNN) classifier is used to infer the topic of a segment embedding.

Fig. 3 illustrates the training and inference process of EMD. The left side shows how segments from day 0 and 1 training EEGs are used to train the embedding model. EEGs for only two pages are shown. The trained embedding model maps EEG segments to the embedding space. Using training data from multiple days helps the embedding model learn what is invariant about the recall of a topic over time.

Fig. 3 (right) shows how labels are inferred for day 1 test EEGs. The segments of all day 0 test EEGs are mapped to the embedding space using the trained embedding model. Segments from the same EEG tend to map to nearby points in the embedding space, even though the test EEGs were not used in training. A KNN classifier is trained on the embeddings of the day 0 test EEGs, and then used to infer the label of each segment of every unlabeled, day 1 EEG. (See (Bruns, Haidar, and Rubino 2023) for a discussion of alternatives to KNN for downstream classifier.)

A key property of this system is that it can predict a topic that the neural embedding model has never been trained on. This is in contrast to classical classifiers that can only infer classes seen in the training data. However, a topic can only be predicted from a day 1 test EEG if a day 0 test EEG of the same topic is available.

Related Work

In this section, we survey some existing applications of neural decoding. Previous studies demonstrate non-invasive brain recordings can be used to reconstruct music and stories experienced during fMRI recordings (Denk et al. 2023; Tang et al. 2023). Although fMRI is an effective tool for neural decoding, its large size and high cost limit its practicality in real-world applications.

We are particularly interested in more portable and inexpensive applications. For example, in education, neural decoding with EEG has been used to classify student engagement levels during presentations, providing insights into learning processes and personalized educational intervention (Poulsen et al. 2017).

Within assistive technology, EEG-based neural decoding has made significant advancements. Recent systems have been used to control prosthetics and facilitate communication (Limchesing et al. 2021). The integration of neural decoding in healthcare promises new methods of rehabilitation and quality-of-life improvements for users (Orban et al. 2022). Additional assistive technologies include real-time epileptic seizure detection and personalized sleep recommendation (Olokodana et al. 2021; Ghosh et al. 2023).

Non-medical, consumer-driven applications include interfacing with video games (Vasiljevic and De Miranda 2020; Putri et al. 2019), predicting consumer preferences of e-commerce products (Amin et al. 2020), smart home control (Qin et al. 2020), and even sports performance (Slutter, Thammasan, and Poel 2021). We are unaware of applications that use EEG data for information retrieval.

Neural Information Retrieval

In this section we define the application of neural information retrieval. The application functions as follows. A user finds a document of interest. To *index* the document, the user briefly recollects the content of the document while an EEG is recorded. The EEG is stored along with a link to the document. Later, to *search* for the document, the user recalls the content of the document while an EEG is again recorded. The application then returns an ordered list of document links, ordered by estimated relevance. Only previously-indexed documents appear in the list.

Any kind of document could be used, provided it has a unique identifier, such as a URL. The document need not be a text document. It could be an image, an audio recording, or a video recording, for example. Further, the application could support multiple types of documents. We generically refer to a unique document identifier as a “link”.

Before using the application, the user would need to personalize the application through a training process. In this process, the user would be asked to record EEG data for a collection of short documents. For each document, the user would first read (or otherwise study) the document, then record EEG data while recalling it.

The design of the EMD classifier (previously described in the section on neural memory decoding) can be adapted for our proposed application. We now describe the training, indexing, and search operations of the application, highlighting differences with the EMD classifier.

Fig. 4 (left) shows how the application is trained. A collection of training EEGs are used to train the embedding model. In contrast to the training of the EMD classifier, shown on the left of Fig. 3, the training EEGs for a document are not necessarily recorded on two consecutive days. For each document, a user could provide a single EEG or multiple EEGs, recorded over any combination of days.

Fig. 4 (middle) shows how a document is indexed. Each segment from the EEG for the document is mapped to the embedding space using the trained embedding model. No retraining of the embedding model is performed. In comparison to the EMD system, the mapping of segments to points in the embedding space is precomputed; not performed at inference time.

Fig. 4 (right) shows how search is performed. The segments of the unlabeled EEG are mapped to the embedding space, and then a KNN classifier infers a label for each of the segments. In comparison to Fig. 3, inference for a document happens any time after it is indexed.

With EMD, training data for a document consisted of a day 0 EEG and a day 1 EEG. In our application, training data for a document can involve EEG data from one or more days, recorded at any time. Another change to EMD in adapting it to our proposed application is retraining of the system over time. With EMD, a fixed set of training documents are used. In an information retrieval system, the number of indexed documents will grow over time, making classification more difficult. This can be compensated for by increasing the size of the training set used to train the neural embedding model. Approaches to retraining over time are discussed in the later section on deployment.

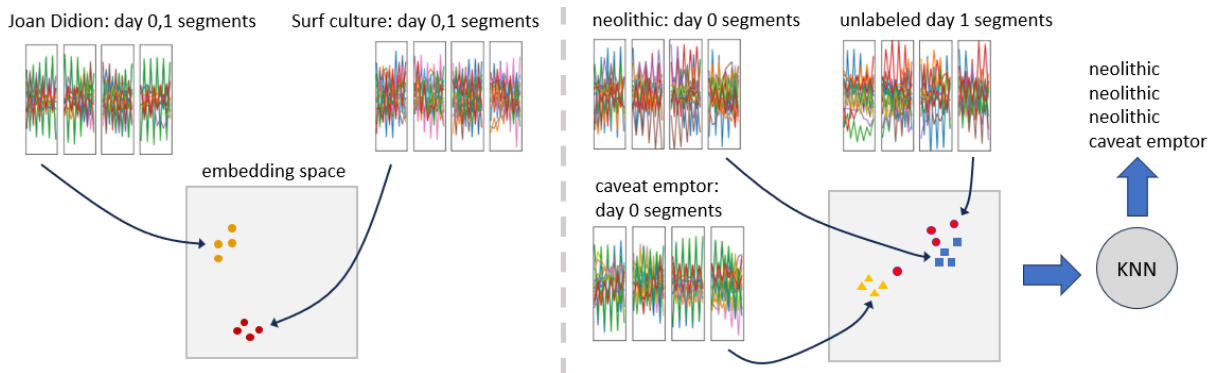


Figure 3: Training (left) and inference (right) in the system of (Bruns, Haidar, and Rubino 2023).

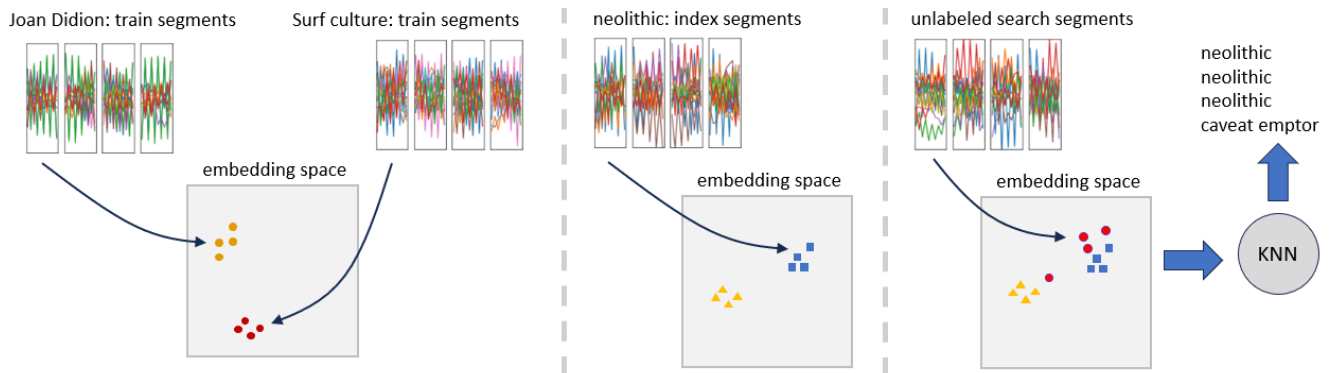


Figure 4: Training (left), indexing a document (middle), and searching for a document (right) in neural information retrieval. Web pages are used as example documents in the figure, but the application could work with any kind of document that has a unique identifier.

Experimental Results

The EMD system described in (Bruns, Haidar, and Rubino 2023) was created to show the feasibility of neural memory decoding with EEG data, not to serve as part of a software application. In preparation for discussion of deployment issues, we present here the results of experiments to test whether our application is practical. Our tests adapt EMD’s code and data, which were provided with the authors’ permissions.

Time spent in EEG recording. Users would not want to spend too much time recording EEG data. In neural information retrieval, an EEG is recorded for each training document, whenever a document is indexed, and whenever a search is performed.

In (Bruns, Haidar, and Rubino 2023), 78 documents were used for training, and all EEGs were 75 seconds long, which amounts to about 850 segments per document. Thus, over 3 hours of EEG recording was performed for system training. Row 1 of Table 1 shows the performance of this baseline system.

To see how the EMD system performs under a more practical scenario, we tested system accuracy with 40 training topics, using 20 segments per training topic, and 10 segments for both indexing and search. For training, following

the process of (Bruns, Haidar, and Rubino 2023), both day 0 and day 1 EEG data was used for each document. The data set for the experiment by using a random subset of the training documents, and a random subset of the segments associated with each document.

A single segment is about 0.8 seconds in length, so under this more practical regime about 21 minutes of EEG recording would be needed for training, and about 8 seconds of EEG recording would be needed for each indexing and search operation.

Row 2 of Table 1 shows the EMD system accuracy using this “small data size” regime. Repeated random subsampling validation was used to compute the accuracy values; each row of the table is the result of at least 20 random splits of the topics into training and test topics. A 95% confidence interval obtained using bootstrapping is shown after the top-1 and top-3 accuracy values. The table shows that the small data regime results in a significant drop in top-1 accuracy, from about 75% to about 48%, and a smaller drop in top-3 accuracy, from about 94% to about 81%. The test set always contains 25 topics, so the top-1 accuracy should be compared to a baseline (obtained by random guessing) of 4%. No hyperparameter tuning was performed after reducing the data size, so the values in the table are conservative.

row	data size	model size	EEG channels	mean top-1	mean top-3
1	full	655k	14	0.750 [0.71, 0.78]	0.944 [0.93, 0.96]
2	small	655k	14	0.476 [0.44, 0.51]	0.812 [0.78, 0.84]
3	full	39k	14	0.690 [0.66, 0.72]	0.918 [0.90, 0.94]
4	full	655k	4	0.764 [0.73, 0.80]	0.954 [0.93, 0.97]
5	small	39k	14	0.472 [0.43, 0.51]	0.808 [0.77, 0.85]
6	small	655k	4	0.512 [0.48, 0.54]	0.884 [0.85, 0.92]
7	full	39k	4	0.726 [0.69, 0.76]	0.922 [0.90, 0.94]
8	small	39k	4	0.500 [0.46, 0.55]	0.852 [0.82, 0.89]

Table 1: Top-1 and top-3 accuracy by system data size, number of model parameters, and number of EEG channels. In the data size column, “full” means 78 training documents and 75 second EEG recording duration for all EEGs, and “small” means 40 training documents, with 16 second EEG recording duration in training, and 8 second duration in indexing and search.

Size of the neural embedding model. Another practical matter is the memory footprint of the embedding model and the CPU work needed in inference. These factors can limit the deployment options for the application. The neural embedding model of EMD has about 655k parameters.

To see the impact of model size, we tested system accuracy using an embedding model of about 39k parameters (about 6% of the size of the model of EMD). The size reduction was achieved by reducing the number of convolutional layers from 4 to 3, reducing the number of dense layers from 2 to 1, reducing the number of convolutional filters from 256 per layer to 64 per layer, and reducing the embedding size from 32 to 24.

Rows 1 and 3 of Table 1 show the result of changing only the size of the embedding model. The mean top-1 accuracy drops from about 75% to about 69%, and the top-3 accuracy drops from about 94% to about 92%. Rows 1 and 5 of the table show the result of changing both the data and model sizes. The accuracy values are similar to those on row 2; reducing the data size has little impact once the model size is reduced.

Number of EEG channels. Fig. 6, left shows the large EEG headset, containing 16 electrodes, that was used to collect EEG in (Bruns, Haidar, and Rubino 2023). Fig. 6, right shows a more comfortable, consumer-grade headset with only 4 electrodes.

The neuroscience literature suggests that the prefrontal cortex plays a large role in the formation and recall of semantic memory (Devlin et al. 2002; Gabrieli, Poldrack, and Desmond 1998). We therefore tested the accuracy of the system using EEG data from only the electrodes in locations FP1, FP2, F3, and F4 of the international 10-20 system (see Fig. 5).

Rows 1 and 4 of of Table 1 show the result of reducing the number of EEG channels from 14 to 4. There is no significant change in top-1 or top-3 accuracy as a result of this change, at least when no change is made to data size or model size.

Rows 1 and 8 of Table 1 show the result of reducing the data size, model size, and number of channels. Top-1 accuracy drops from about 75% to about 50%, and top-3 accuracy drops from about 94% to about 85%. The drop in top-3 accuracy is significant, but represents a top-3 accuracy reduction of only about 10%.




Figure 5: Electrode locations in the 10-20 system.

Deployment and Data Privacy

A consumer application or service for neural information retrieval should satisfy requirements that include:

1. Users’ EEG data must be kept private.
2. The EEG recording equipment should be affordable and comfortable.
3. The application should be usable in a range of environments.
4. The time spent in recording EEG data should not be a burden.
5. The application should be able to be run with acceptable performance on typical consumer laptops and, ideally, also smartphones.
6. The performance of the system should not degrade as more and more documents are indexed over time.

The first requirement is important because the health information that could be derived from EEG data is unknown today. We prefer a deployment model in which EEG data is never shared with a service provider.

Regarding the second requirement, the EEG sensing hardware used in (Bruns, Haidar, and Rubino 2023) costs about US\$ 3000, which includes a 16-channel headset, electrodes, and interface hardware (see Figure 6, left). More comfortable and lower-cost consumer-grade headsets are also available. For example, the one shown in Figure 6 (right) sells

for about US\$ 400, including all hardware. It has four EEG electrodes, two of which are near the prefrontal cortex. Results in the neuroscience literature link the prefrontal cortex with semantic memory recall (Gabrieli, Poldrack, and Desmond 1998; Martin and Chao 2001). The importance of these nodes in memory recall is further substantiated by experimental results presented in Table 1, which show that the accuracy of EMD can be matched when only four electrodes are used: the electrodes associated with the prefrontal cortex, and two of the frontal nodes.

A range of lower-cost headsets are assessed in (LaRocco, Le, and Paeng 2020; Pathirana, Asirvatham, and Johar 2018) in the context of drowsiness detection and brain-computer interfaces.

The third requirement, concerning use of the application in a range of environments, is important because EEG is sensitive to noise and is traditionally performed in a clinical environment. We do not know how the accuracy of the application would be affected by noise in the recording environment, or whether the EEG artifacts of this noise could be suppressed by an application. This topic is a subject for future work.

The fourth requirement is addressed by the data size experiments presented in the experimental results section. They show that the top-3 accuracy of the EMD system is about 85% when only 8 seconds of EEG capture is used for indexing and search. We do not have data on the length of time users typically require to bookmark a page or to search for an existing bookmark. However, experience suggests that 8 seconds is comparable to the time needed for these tasks.

The fifth requirement, concerning use of the application on laptops, smartphones, and other mobile devices, is tied to how the application is deployed: as a web service, locally in the browser, or locally as an application. To achieve data privacy, we focus on the latter two options. Suppose the application is run locally on a modern consumer laptop. Does it have the memory, storage, and CPU resources needed for our application to perform well?

To answer this question, we look at the operations of training, indexing, and search. Indexing requires only that a recorded EEG be segmented, preprocessed, and then each segment mapped to an embedding using the trained embedding model. With a small embedding model, such as the one discussed in the experimental results section, these steps can easily be handled on a laptop.

Search involves the same steps as indexing, but additionally requires the classification of the embeddings of search segments using an instance-based classifier. Tests with the Scikit-Learn machine learning library, made under the the assumption of 20K topics and 10 segments used in indexing and search, show that the KNN classification step could be performed in less than 0.1s on a 2017 Windows computer, and so could also be performed acceptably fast on a modern laptop.

Training involves the training of the embedding model, which is a convolutional network. Training the model of (Bruns, Haidar, and Rubino 2023) on a 2017 Windows computer, using only the CPU, takes only about 4 minutes. Regarding storage, a 4-channel EEG of length 20 seconds re-



Figure 6: The OPENBCI Mark IV EEG headset (left). The Muse-S EEG headband (right). Image sources: openbci.com and choose muse.com.

quires only about 40kB of storage, or 40MB for 1000 such EEGs – the equivalent of about 20 high quality JPEG images.

Regarding the sixth requirement, simple occasional re-training of the embedding model using EEGs acquired in the indexing process may be sufficient. Alternatively, approaches like incremental learning, in which a model is improved using new training data, without access to earlier training data, may also be an option. Particularly relevant to our application is few-shot class incremental learning (Tao et al. 2020).

Conclusions

We have defined the application of neural information retrieval, described the design of its machine learning component, and presented experimental results related to its deployment.

More experimentation is needed to address questions of the practicality of the application. How much will performance of the application vary from person to person? How will the performance change as thousands of documents are indexed? Will the system perform well when the time between indexing a document and searching for it is weeks, months, or even years?

We illustrated neural information retrieval with a text retrieval application, but neural information retrieval is not limited to documents consisting of text or to semantic memories of concepts. For example, in a music retrieval application, a user could retrieve an audio recording by recalling the sound of a passage in a piece of music. We believe a wide range of useful applications will be discovered.

Acknowledgments

The work of Michael Haidar, and the equipment used in this project, was partly funded by the Undergraduate Research Opportunities Center at California State University, Monterey Bay, the Koret Foundation, and the Office of the Director of the National Institutes of Health under grant number R25MD010391. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- Amin, C. R.; Hasin, M. F.; Leon, T. S.; Aurko, A. B.; Tamanna, T.; Rahman, M. A.; and Parvez, M. Z. 2020. Consumer behavior analysis using EEG Signals for neuromarketing application. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2061–2066.
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828.
- Bird, J. J.; Ekart, A.; Buckingham, C. D.; and Faria, D. R. 2019. Mental emotional sentiment classification with an EEG-based brain-machine interface. In *Proceedings of DISP '19*.
- Bruns, G.; Haidar, M.; and Rubino, F. 2023. Neural memory decoding with EEG data and representation learning. arXiv:2307.13181.
- Chadwick, M. J.; Hassabis, D.; Weiskopf, N.; and Maguire, E. A. 2010. Decoding individual episodic memory traces in the human hippocampus. *Current Biology*, 20(6): 544–547.
- Denk, T. I.; Takagi, Y.; Matsuyama, T.; Agostinelli, A.; Nakai, T.; Frank, C.; and Nishimoto, S. 2023. Brain2Music: Reconstructing music from human brain activity. arXiv:2307.11078.
- Devlin, J. T.; Russell, R. P.; Davis, M. H.; Price, C. J.; Moss, H. E.; Fadili, M. J.; and Tyler, L. K. 2002. Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. *Neuropsychologia*, 40(1): 54–75.
- Gabrieli, J. D. E.; Poldrack, R. A.; and Desmond, J. E. 1998. The role of left prefrontal cortex in language and memory. *Proceedings of the National Academy of Sciences*, 95(3): 906–913.
- Ghosh, K.; Nanda, S.; Hurt, R. T.; Schroeder, D. R.; West, C. P.; Fischer, K. M.; Bauer, B. A.; Fokken, S. C.; Ganesh, R.; Hanson, J. L.; et al. 2023. Mindfulness using a wearable brain sensing device for health care professionals during a pandemic: A pilot program. *Journal of Primary Care & Community Health*, 14: 21501319231162308.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- LaRocco, J.; Le, M. D.; and Paeng, D.-G. 2020. A systemic review of available low-cost EEG headsets used for drowsiness detection. *Frontiers in Neuroinformatics*, 14.
- Limchesing, T. J. C.; Chua, A. E. C.; Shi, C. J. L.; Baldovino, R. G.; Munsayac, F. E. T.; and Bugtai, N. T. 2021. A review on recent applications of EEG-based BCI in wheelchairs and other assistive devices. In *Proceedings of HNICEM 2021*, 1–6. IEEE.
- Martin, A.; and Chao, L. L. 2001. Semantic memory and the brain: structure and processes. *Current Opinion in Neurobiology*, 11(2): 194–201.
- Ofen, N.; Whitfield-Gabrieli, S.; Chai, X. J.; Schwarzlose, R. F.; and Gabrieli, J. D. 2017. Neural correlates of deception: Lying about past events and personal beliefs. *Social cognitive and affective neuroscience*, 12(1): 116–127.
- Olokodana, I. L.; Mohanty, S. P.; Kougianos, E.; and Sherratt, R. S. 2021. EZcap: A novel wearable for real-time automated seizure detection from EEG signals. *IEEE Transactions on Consumer Electronics*, 67(2): 166–175.
- Orban, M.; Elsamanty, M.; Guo, K.; Zhang, S.; and Yang, H. 2022. A review of brain activity and EEG-based brain-computer interfaces for rehabilitation application. *Bioengineering*, 9(12).
- Pathirana, S.; Asirvatham, D.; and Johar, G. 2018. A critical evaluation on low-cost consumer-grade electroencephalographic devices. In *Proceedings of ICBAPS '18*, 160–165.
- Polyn, S.; Natu, V.; Cohen, J.; and Norman, K. 2006. Category-specific cortical activity precedes retrieval memory search. *Science (New York, N.Y.)*, 310: 1963–6.
- Poulsen, A. T.; Kamronn, S.; Dmochowski, J.; Parra, L. C.; and Hansen, L. K. 2017. EEG in the classroom: Synchronised neural recordings during video presentation. *Scientific reports*, 7(1): 43916.
- Putri, F. D.; Ding, H.; Garcia, A.; and Vuckovic, A. 2019. Towards successful multi-user brain-computer interface (BCI) gaming: Analysis of the EEG signatures and connectivity. In *CHIRA*, 59–65.
- Qin, L. Y.; Nasir, N. M.; Huq, M. S.; Ibrahim, B.; Narudin, S. K.; Alias, N. A.; and Ab Ghani, M. A. 2020. Smart home control for disabled using brain computer interface. *International Journal of Integrated Engineering*, 12(4): 74–82.
- Rissman, J.; Greely, H. T.; and Wagner, A. D. 2010. Detecting individual memories through the neural decoding of memory states and past experience. *Proceedings of the National Academy of Sciences*, 107(21): 9849–9854.
- Ritter, P.; and Villringer, A. 2006. Simultaneous EEG–fMRI. *Neuroscience & Biobehavioral Reviews*, 30(6): 823–838.
- Slutter, M. W. J.; Thammasan, N.; and Poel, M. 2021. Exploring the brain activity related to missing penalty kicks: An fNIRS study. *Frontiers in Computer Science*, 3.
- Tang, J.; LeBel, A.; Jain, S.; and Huth, A. G. 2023. Semantic reconstruction of continuous language from non-invasive brain recordings. *Nature Neuroscience*, 1–9.
- Tao, X.; Hong, X.; Chang, X.; Dong, S.; Wei, X.; and Gong, Y. 2020. Few-shot class-incremental learning. In *Proceedings of CVPR 2020*, 12183–12192.
- Tulving, E. 1972. Episodic and semantic memory. In *Organization of memory*, 381–403. Academic Press.
- Vargas, R.; and Just, M. A. 2020. Neural representations of abstract concepts: identifying underlying neurosemantic dimensions. *Cerebral Cortex*, 30(4): 2157–2166.
- Vasiljevic, G. A. M.; and De Miranda, L. C. 2020. Brain-computer interface games based on consumer-grade EEG devices: A systematic literature review. *International Journal of Human-Computer Interaction*, 36(2): 105–142.
- Winocur, G.; and Moscovitch, M. 2011. Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5): 766–780.